

Heroic Reliability Improvement in Manned Space Systems

Harry W. Jones¹

NASA Ames Research Center, Moffett Field, CA, 94035-0001

System reliability can be significantly improved by a strong continued effort to identify and remove all the causes of actual failures. Newly designed systems often have unexpected high failure rates which can be reduced by successive design improvements until the final operational system has an acceptable failure rate. There are many causes of failures and many ways to remove them. New systems may have poor specifications, design errors, or mistaken operations concepts. Correcting unexpected problems as they occur can produce large early gains in reliability. Improved technology in materials, components, and design approaches can increase reliability. The reliability growth is achieved by repeatedly operating the system until it fails, identifying the failure cause, and fixing the problem. The failure rate reduction that can be obtained depends on the number and the failure rates of the correctable failures. Under the strong assumption that the failure causes can be removed, the decline in overall failure rate can be predicted. If a failure occurs at the rate of λ per unit time, the expected time before the failure occurs and can be corrected is $1/\lambda$, the Mean Time Before Failure (MTBF). Finding and fixing a less frequent failure with the rate of $\lambda/2$ per unit time requires twice as long, time of $1/(2\lambda)$. Cutting the failure rate in half requires doubling the test and redesign time and finding and eliminating the failure causes. Reducing the failure rate significantly requires a heroic reliability improvement effort.

Nomenclature

<i>CCF</i>	=	Common Cause Failures
<i>CDRA</i>	=	Carbon Dioxide Removal Assembly
<i>ECLS</i>	=	Environmental Control and Life Support
<i>ISS</i>	=	International Space Station
<i>MTBF</i>	=	Mean Time Before Failure
<i>OGA</i>	=	Oxygen Generation Assembly

I. Introduction

RELIABILITY growth requires heroic effort to detect and prevent failures. This paper describes the reliability life cycle, presents mathematical models of reliability growth, analyzes published reliability growth data, and studies the reliability of manned space systems including life support. If a system is operated over time, failures occur, and if the observed failure modes are eliminated, reliability growth will occur. The reduction in failure rate depends on the initial failure numbers and rates. Achieving very low final failure rates requires a heroic reliability effort because a long test time and continuing effort is needed for failures with low rates to occur and be removed.

If all the detected failure modes are eliminated, reliability grows at the same rate as the failures occur. Reliability growth can be tracked as the expected decline in the current failure rate, and this depends directly on the failure rates of the remaining undetected failure modes. This process includes the detection and repair of common cause failures, which can be due to easily correctable design mistakes. Hardware failures may be due to random events or unpreventable degradation, but some failures causes are preventable. All software failures or bugs are design errors, essentially common cause failures because they affect all copies of the software, and they are usually fixed. Insight into the reliability growth mechanism can be gained by considering different numbers of failure modes with specific arrays of failure rates.

Some detected failure modes may not be corrected, especially if they are considered random and unpredictable and occur at a low rate. Some failures may not be preventable with current technology. The more usual reliability assumption is that most failures are random and independent, they occur at a low rate, they can be repaired using a

¹ Systems Engineer, Bioengineering Branch, Mail Stop N239-8.

small stock of spare parts, and they do not indicate a need for redesign. Using two classes of failures, correctable and not correctable or random, allows a better model for actual reliability growth data than the standard reliability growth model that assumes reliability growth continues indefinitely.

II. Reliability growth background

Failure rates usually vary over time. A decreasing failure rate, called reliability growth, is often observed. Reliability growth models can track and attempt to predict a declining failure rate.

A. The system failure rate changes over time according to the “bathtub curve”

The failure rate is the number of failures per unit time. A system’s changing failure rate over time often follows the “bathtub curve.” The failure rate first decreases with time, then remains constant during the system’s useful life, and finally increases due to component wear out. The initial high “infant mortality” failure rate is due to burn-in, to failure of defective components, to detection and correction of design faults, and to other improvements in design and operations. The failures during useful life are usually assumed to be random events caused by unpredictable internal degradation. The failure rate increase at end-of-life can be caused by mechanical wear or aging related to chemical or thermal activity. The bathtub curve is shown in Figure 1.




Figure 1. The bathtub failure rate curve, arbitrary time units.

More than two-thirds of all systems show infant mortality and then a constant failure rate, but then have no final aging period. (Hansen, 2001) Some systems such as the space shuttle have long continued reliability growth, probably because of extensive maintenance and refurbishment. (Shishko, 1995)

B. The Duane-Crow reliability growth model

Commonly used reliability growth models assume that the failure rate decreases over time and that it declines as a function of time to the minus (roughly) one-half power. The exact downward slope is found using a log-log graph of failure rate versus time.

The failure rate, λ , is the number of times a component or system is expected to fail per unit time, given that it is currently still operating. It is usually assumed that the failure rate, λ , is constant during the useful life of a system. The Mean Time Before Failure (MTBF) is the inverse of the failure rate, λ .

$$\text{MTBF} = 1/\lambda$$

Reliability growth, a decreasing failure rate over time, occurs when design improvements are made and failure modes removed. When the failure rate varies with time, the instantaneous failure rate is written $\lambda(t)$.

Duane observed in 1964 that if $N(t)$ is the number of failures occurring until time t , a plot of the cumulative failure rate, $N(t)/t$, versus the cumulative test time, t , often closely follows a straight line when plotted on log-log graph paper. The observed relation is

$$N(t)/t = k t^{-\alpha}$$

The parameter α , the growth rate, is the downward slope of $N(t)/t$ versus t . Measured values of α , the growth rate parameter, usually vary from 0.2 to 0.6. (Yamada and Osaki, 1983) (MIL-HDBK-189C, 2011)

Crow in the 1970's provided a theoretical basis for the Duane model. He assumed that the failures of a system during development testing occur according to a non-homogeneous (time-varying) Poisson process with a power law mean value function, $m(t)$. The mean number of failures over time is assumed to be

$$m(t) = k t^{\beta}$$

where β is between zero and one.

The instantaneous failure rate is the time derivative of the number of failures.

$$\lambda(t) = d[m(t)]/dt = k \beta t^{\beta-1}$$

This is known as the Weibull distribution failure rate, although the full Weibull distribution is more complex. The mathematically expected cumulative failure rate is given by

$$\text{Expected } [N(t)/t] = m(t)/t = k t^{\beta-1}$$

The Crow and Duane reliability growth models are equivalent, with the Duane α equal to Crow's $1 - \beta$. The parameter k is the same in both. The parameter β is the ratio of the current instantaneous failure rate, $\lambda(t)$, to the average cumulative failure rate, $m(t)/t$.

$$\beta = \lambda(t)/[m(t)/t] = k \beta t^{\beta-1} / k t^{\beta-1}$$

The typical β of 0.4 to 0.8 corresponds to a decreasing failure rate and positive reliability growth. (Yamada and Osaki, 1983) (MIL-HDBK-189C, 2011)

The reliability growth parameters can be estimated from failure time data. Suppose that N failures are observed during the test time $(0, T)$, and that they occur sequentially at times s_1, s_2, \dots, s_N . The maximum likelihood estimate of β is

$$\beta^* = N / \sum \ln (T/s_i)$$

where \ln is the natural logarithm and the summation \sum is over $i = 1$ to N . The maximum likelihood estimate of k is

$$k^* = N / T^{\beta^*}$$

(Yamada and Osaki, 1983) (MIL-HDBK-189C, 2011)

C. Applying the Duane-Crow reliability growth model

Crow used a data set of 56 failures occurring over 400 hours to illustrate the reliability growth model. (MIL-HDBK-189C, 2011). A graphical Duane model fit to this data gives

$$\text{Duane } N(t)/t = 0.640 t^{-0.283}$$




Figure 2. The failure rate $N(t)/t$ data and the Duane line and Crow model fits.

Figure 2 shows the cumulative failure rate data $N(t)/t$ plotted versus time, t , in a log-log graph, along with the Duane line fit and Crow model fit.

The downward slope of the Duane line fit shows significant reliability growth, but this is due to a high early failure rate, infant mortality. If the first two failures are ignored, the Duane line downward slope is much shallower, corresponding to $\alpha = 0.171$. (Jones, 2014-075)

Crow's analysis of this data set using the above formulas found $k = 0.217$ and $\beta = 0.927$.

$$\text{Crow } [N(t)/t] = k t^{\beta-1} = 0.217 t^{-0.073}$$

The β corresponds to an $\alpha = 1 - \beta = 0.073$, which is much less than the $\alpha = 0.283$ found using the Duane graphical method. The Crow model is much less influenced by the early infant mortality data and gives a more pessimistic projection of future reliability growth. (Jones, 2014-075) Nevertheless, reliability growth analysts claim that, "While growth is small, hypothesis testing indicates it is significantly different from 0. Thus growth is occurring and the failure intensity (failure rate) is decreasing." (MIL-HDBK-189C, 2011) This seems an exaggeration.

D. The example data does not show continuing reliability growth

It is obvious that reliability growth has ceased long before $t = 400$ in this data set. Figure 3 shows the cumulative failure rate, $N(t)/t$ plotted versus time, t , but in a linear rather than log-log graph. A two part curve is fit is also shown, rather than a single curve fit as usually used.




Figure 3. Cumulative failure rate $N(t)/t$ and a two part fit in a linear graph.

Each $N(t)/t$ data point is a failure occurrence, so if the next failure is long delayed, the $N(t)/t$ cumulative failure rate declines. After the high initial failure rate is smoothed out over time, the cumulative failure rate in Figure 3 becomes constant. A flat line fits the data points from time 100 to 400. A Duane-Crow form equation with $N(t)/t = 1.11 t^{-0.5}$ fits the data from time 0 to 100. A log-log plot of the data in Figure 3 is shown in Figure 4.




Figure 4. Cumulative failure rate $N(t)/t$ and a two part fit in a log-log graph.

The $\alpha = 0.5$ is a typical expected slope for reliability growth. This data set that has been used in the literature to explain reliability growth clearly does not show continuous reliability growth. As commonly illustrated by the bathtub curve, an early period of strong reliability growth with a constantly declining failure rate is followed by a longer period with a lower constant failure rate. The green flat line after $t = 100$ is obscured by the data points and the log plot exaggerates the data variability. The Duane approach above of fitting a curve with $N(t)/t = k t^{-\alpha}$ to the full data set produced $\alpha = 0.283$, which is roughly the average of the early $\alpha = 0.5$ reliability growth and the later $\alpha = 0$ constant failure rate.

E. The Duane-Crow model mistakenly assumes that reliability growth continues

The Duane-Crow model fits the failure rate data with a single line. It is an obvious and well known fact, commonly illustrated by the “bathtub curve,” that failure rates usually decline strongly during an initial period of “infant mortality,” and then tend to be constant during a longer operational phase. Data from the initial period of upgrades shows significant reliability growth usually caused by strong efforts to test the system, discover its failures, and eliminate them. Using a Duane-Crow model line fit can exaggerate the long term reliability growth potential if it uses only the early data.

Continuing the reliability upgrade process indefinitely or until a reliability target is reached requires heroic effort. Adding more and more data from the later operational period gives a false picture of continuing reliability growth occurring at an ever decreasing rate. If the test and fix process is terminated and the system put into operation, the best predictor of the future failure rate would be the failure rate at the end of the reliability growth effort, assuming no “end-of-life” increase in failure rate occurs.

Clearly the two phases, initial improvement and continued operation, have different failure rate behavior and require different analysis and modeling. Combining reliability growth and constant reliability will be considered in the next section.

III. Reliability growth modeling based on assumed discrete failure rates

This section develops reliability growth models considering different correctable failure rates.

A. Reliability growth is due to fixing the most frequent (highest rate) failure modes

Reliability growth will occur in a system if its failure causes can be removed by redesign or otherwise prevented. A mature system that has been well designed, thoroughly tested, and operated successful for a long history will have occasional failures, but these are usually considered random, unpredictable, and unpreventable. In this case, a failure probably does not lead to a redesign since there is little need to improve reliability. But an all new design may have an unacceptably high failure rate, “infant mortality,” due to design oversights, specification errors, improper operation, etc. Some failure modes are considered common cause failures or CCFs rather than mistakes. Curing any

failure mode increases reliability. Reliability growth is expected in early testing and may be considered necessary if the initial reliability is too low.

A design with no existing curable problems is preferable, but design errors and improvement opportunities often occur in new systems. It is necessary to identify and correct problems as far as possible. To explain reliability growth, it is necessary to consider the number of curable failure modes and the process of finding them. As an example, suppose there is an array of standard, well designed LEDs with high reliability and long life, say 100,000 hours. They can be assumed to have a constant failure rate of 1 in 100,000 hours due to random unpredictable events. Suppose that there are many large LED arrays, so that some LEDs are always failing and being replaced, and also suppose that after some time it is noticed that some of the new replacement LEDs are failing very rapidly, within a few months, on average 1 in 1,000 hours. The failure investigation shows a design change or an assembly machine failure or a material contamination and the problem is fixed. In order to be fixed, a failure mode must occur, and the probability of its occurrence over time depends on its failure rate. The higher the failure rate, λ , the sooner the failure occurs. The MTBF = $1/\lambda$. The overall failure rate of a system is the sum of the component failure rates. The failure modes with the highest failure rates will occur first. If these observed high rate failure modes are then cured, the failure rate declines rapidly and reliability growth occurs.

B. Reliability growth and failure rate reduction for a single correctable failure

Consider a new design with multiple failure modes, including independent random failures and correctable failures. Reliability growth depends on the failure rates of the correctable failures, although random failures can have a masking effect. Suppose that the combined failure rate of all the non-correctable failure modes is λ_{non} . Suppose that the combined failure rates of all the correctable failure modes is λ_{cor} . The total initial failure rate is $\lambda_{total} = \lambda_{non} + \lambda_{cor}$. The minimum final failure rate after all possible reliability growth is $\lambda_{min} = \lambda_{non}$.

How does reliability growth occur over time? The failure process is probabilistic. Suppose that there is only one correctable failure with failure rate λ_{cor} . The probability of this failure having occurred over time is $1 - \exp(-\lambda_{cor} t)$, where t is time. The expectation of the failure occurring and being removed gradually increases from zero to one, so the expected failure rate decreases from $\lambda_{total} = \lambda_{non} + \lambda_{cor}$ to $\lambda_{min} = \lambda_{non}$. This is a process of exponential decay, $\lambda(t) = \lambda_{non} + \lambda_{cor} * \exp(-\lambda_{cor} t)$. Figure 5 shows the expected failure rate decline for a single correctable failure. Actually, the single correctable failure will occur at one point in time, and when it is corrected $\lambda(t)$ drops from λ_{total} to $\lambda_{min} = \lambda_{non}$.

C. Failure rate reduction for several correctable failures

Suppose there are several correctable failures. The total failure rate is the sum of the individual failure rates. If there are N of them, then $\lambda_{cor} = \sum_{i=1}^N \lambda_{cor_i}$. Each of the failures occurs according to its own failure rate λ_{cor_i} . Each has its contribution to the total correctable failure rate that declines with time, $\lambda_{cor_i}(t) = \lambda_{cor_i} * \exp(-\lambda_{cor_i} t)$. The total correctable failure rate declines with time, so that $\lambda_{cor}(t) = \sum_{i=1}^N \lambda_{cor_i} * \exp(-\lambda_{cor_i} t)$.

The summation for $\lambda_{cor}(t)$ does not have a simple solution unless all the λ_{cor_i} failure rates are equal. Since this is the worst case, giving the slowest reliability growth, it is considered next. For all i , $\lambda_{cor_i} = \lambda_{cor}/N$, the total correctable failure rate divided by N . Then $\lambda_{cor}(t) = \sum_{i=1}^N \lambda_{cor}/N * \exp[-(\lambda_{cor_i}/N) t] = \lambda_{cor} * \exp[-(\lambda_{cor}/N) t]$.

Figure 5 also shows the expected failure rate decline for ten correctable failures with the same failure rate. If instead of one correctable failure with the rate λ_{cor} , we have N failures each with the same failure rate, λ_{cor}/N , the time for reliability growth is increased by N . It takes N times longer to achieve the same reliability growth. Reliability growth is more rapid for fewer, higher failure rate correctable failures.

Consider an example. If there is only one correctable failure with a failure rate of 1 in 100 time units, it will probably occur and be fixed in 100 time units. If there are 10 correctable failures each with failure rate of 1 in 1,000 time units, the overall correctable failure rate remains 1 in 100 time units. The first one will probably occur and be fixed in 100 time units, but that will leave nine undetected and unfixed. It will require 1,000 time units before we can expect to find and fix all 10 correctable failures. The time for reliability growth is increased by a factor of 10.




Figure 5. Reliability growth due to discovery and removal of either one or ten correctable failures.

Having ten instead of one correctable failure with the same total failure rate stretches out the expected time for reliability growth by a factor of ten. If instead of one highly probable failure mode, there are ten failure modes with one-tenth the original probability, it will take ten times longer to find and fix them. The potential amount of reliability growth equals the total correctable failure rate. The time to achieve reliability growth depends on the individual MTBFs of the individual failure modes.

D. A simple failure rate bound defining the minimum heroic reliability growth

Figure 5 shows how the average failure rate declines over time as correctable failure modes are found and eliminated. A simple reliability growth bound can be derived without knowing the actual correctable failure rates. (Bishop and Bloomfield, 1996)

The average expected failure rate for each correctable failure declines with time, $\lambda_{cor_i}(t) = \lambda_{cor_i} * \exp(-\lambda_{cor_i} t)$. At any given time, there is some initial λ_{cor_i} that maximizes its current average expected failure rate. Taking the derivative of $\lambda_{cor_i}(t)$ with respect to λ_{cor_i} and setting it to zero, it can be shown that the maximum value of $\lambda_{cor_i}(t)$ occurs at the time when t equals $1/\lambda_{cor_i}$, the MTBF. A failure mode with a given failure rate is more likely to occur near the time equal to its MTBF. The maximum value of $\lambda_{cor_i}(t)$ over t is a bound on the current expected remaining failure rate. The bound, found by substituting t equals $1/\lambda_{cor_i}$, is $\lambda_{cor_i}(t) \leq 1/(e t)$. (Bishop and Bloomfield, 1996) Figure 6 plots the bound $1/(e t)$ and the individual expected failure rates over time $\lambda_{cor_i} * \exp(-\lambda_{cor_i} t)$, for $\lambda_{cor_i} = 0.1, 0.01, 0.001, 0.0001, \text{ and } 0.00001$.




Figure 6. The expected failure rate is always lower than the bound $1/(e t)$, regardless of the original failure rate. (Bishop and Bloomfield, 1996)

For each individual failure mode, the expected failure rate declines exponentially with time. The current expected failure rate is always less than, better than, the bound $1/(e t)$, regardless of the original failure rate. The assumptions are that the failure rates are constant and independent and that a failure is immediately corrected, without introducing a new failure mode. The bound directly decreases with increasing test and redesign time. The bound is surprising because it proves reliability growth must occur under the assumptions and because it allows the maximum future failure rate to be predicted from the current failure rate. The bound does not tell us when or even if any particular failure mode will occur, but it sets the maximum expected failure rate. The existence of the bound shows that reliability growth will occur due to detecting and removing correctable failures.

The failure rate bound of $\lambda_{cor_i}(t) \leq 1/(e t)$ is tight only near $t = \text{MTBF}$. The expected failure rate exactly equals the bound when $t = 1/\lambda_{cor_i}$, the MTBF.

E. Failure rate and failure rate bound for several widely different failure rates

If there are N different failure modes, the bound on the total failure rate is $N/(e t)$. If all the N failure modes have the same initial failure rate λ_{cor_i} , the bound is tight at the time equal to the MTBF. However, if the individual failure rates are very different, the total failure rate can be significantly less than the bound. Suppose as above that there are five different failure modes with the initial rates $\lambda_{cor_i} = 0.1, 0.01, 0.001, 0.0001$, and 0.00001 . The sum of the declining expected average failure rates and the bound $N/(e t)$ are shown in Figure 7.




Figure 7. The bound $N/(e t)$ for five failure modes with different initial failure rates.

The sum of the expected failure rates declines as $1/t$ and is always less than the bound $N/(e t) = 5/(e t)$ by a factor of 4. If the λ_{cor_i} are widely different, the total failure rate is much less than the bound and the time to achieve reliability growth stretches to the longest MTBF. Comparing Figures 6 and 7, it can be seen that for widely spaced λ_{cor_i} , only the few failure modes with MTBFs close to the current time contribute substantially to the current $\lambda(t)$.

IV. Reliability growth modeling assuming a continuous gamma failure rate distribution

This section extends the reliability growth model by assuming that the correctable failure rates have a continuous gamma distribution. This provides further insight into the effect of the failure rates on the potential for reliability growth, assuming that the failures are removed once they occur. Fitting failure rate data to a gamma distribution can help explain observed reliability growth.

A. The gamma distribution

It is assumed that the failure rates have the gamma distribution. (Bishop and Bloomfield, 1996)

$$\begin{aligned} \text{Gamma distribution } (\lambda) &= N \text{ gamma } (\lambda, \alpha, \beta) \\ &= N / [\beta \Gamma(\alpha)] (\lambda/\beta)^{\alpha-1} \exp(-\lambda/\beta) \end{aligned}$$

The gamma distribution is shown in Figure 8.




Figure 8. The gamma distribution.

The mean value of the gamma distribution is $\alpha \beta$ and its variance is $\alpha \beta^2$. The gamma distribution becomes an exponential distribution for $\alpha = 1$, as shown in Figure 8. A small mean value $\alpha \beta$, such as 0.5, corresponds to mostly very low failure rates. The small variance of $\alpha \beta^2$ of 0.5 corresponds to a narrow cluster of similar failure rates surrounding the mean of $\alpha \beta$ equal to 5.

B. Heroic expected reliability growth for failure rates with different gamma distributions

The estimated failure rate over time can be derived assuming that the N correctable failure mode rates are distributed according to the gamma distribution. (Bishop and Bloomfield, 1996)

$$\begin{aligned}
 \lambda_{cor}(t) &= \int \text{Distribution}(\lambda) \exp(-\lambda t) \lambda d\lambda \\
 &= \int N/[\beta \Gamma(\alpha)] (\lambda/\beta)^{\alpha-1} \exp(-\lambda/\beta) \exp(-\lambda t) \lambda d\lambda \\
 &= \int N/ \Gamma(\alpha) (\lambda/\beta)^{\alpha} \exp(-1/\beta - t)\lambda d\lambda, \text{ evaluated over } \lambda \text{ from } 0 \text{ to } \infty. \\
 &= \alpha \beta N/(1 + \beta t)^{-(\alpha+1)} \quad (\text{Bishop and Bloomfield, 1996})
 \end{aligned}$$

The gamma distribution becomes the exponential distribution for $\alpha = 1$.

$$\text{Exponential distribution}(\lambda) = N (1/\beta) \exp(-\lambda/\beta)$$

The exponential distribution seems a reasonable failure rate distribution, since it has relatively few components with very high failure rates and many more with lower failure rates. For the exponential distribution the failure rate decline is

$$\lambda_{cor}(t) = \beta N/(1 + \beta t)^{-2}$$

The failure rate declines for gamma distributions of failure rates are shown in Figure 9.




Figure 9. Failure rate declines for gamma distributions of failure rates.

The minimum failure rate decline for any distribution is the bound $1/(e t)$. The failure rate decline for the clustered failures at $\alpha = 50$ and $\beta = 0.1$ touches this bound at the average MTBF of $1/(\alpha \beta) = 0.2$ time units. This is the MTBF corresponding to the average failure rate of $\alpha \beta = 5$. For testing well beyond the MTBF, the failure rate falls rapidly as $t^{-\alpha}$ or t^{-50} . The failure rate decline for the exponentially distributed failures at $\alpha = 1$ and $\beta = 2$ lies below the $1/(e t)$ bound and ultimately declines as t^{-2} . For $\alpha = 0.5$ and $\beta = 1$, there are many low rate failures and the failure rate decline is slower than the other gamma cases.

V. The reliability growth curve, mathematical model, and heroic metrics

The reliability growth description provides an overall graphic picture, a mathematical model of the graph, and an explanation of how the failure rates of correctable failures affect reliability growth. These will be reviewed and used to provide metrics for heroic reliability growth.

A. The reliability growth graph and mathematical model

The plot of the expected failure rate over time is similar to Figure 1 without the end-of-life effect and also to Figures 3 and 5. The expected reliability growth graph is shown in Figure 10.




Figure 10. The reliability growth curve.

The reliability growth curve shown in Figure 10 has two components, a correctable and therefore declining failure rate equal to $0.9 t^{-0.5}$ and a constant random failure rate of 0.1. At time equal to 50 time units, the failure correction process is terminated. A correctable failure rate of 0.13 remains, producing a constant total failure rate of 0.23 after time 50. Continuing to remove the remaining failure modes would have produced the slowly declining failure rate shown.

The mathematical model used to generate the reliability growth curve in Figure 10 is

$$\begin{aligned} \text{Failure rate} &= 0.9 t^{-0.5} + 0.1 \text{ from } t = 0 \text{ to } 50 \\ &= 0.9 (50)^{-0.5} + 0.1 = 0.23 \text{ after } t = 50 \end{aligned}$$

A simple “abcd” mathematical model for reliability growth and failure rate decline is

$$\begin{aligned} \text{Failure rate} &= a t^{-b} + c \text{ from } t = 0 \text{ until the time } t_d \text{ when reliability growth stops and the failure rate} \\ &= c + d \text{ after } t_d \end{aligned}$$

B. The failure rates determine the reliability behavior

The usual reliability model for an operational system corresponds to a constant failure rate, equal to “c” in Figure 10. The overall failure rate “c” is the sum of the individual component failure rates. The failures are assumed to be random and independent and to occur at reasonably low rates. It is assumed that the system can be kept operational by replacing parts that fail. The number of spares required for a component for a particular time period is determined by its failure rate and the acceptable probability of not having a needed spare. The failure rates may be based on component life test data or taken from a database. Sometimes a component has a greater than expected failure rate and a new component or redesign is required. Such failure mode correction is more expected during initial reliability growth phase in Figure 10.

If a system is tested until failures occur, and the failure causes are identified and removed, reliability growth will occur. Under the heroic assumption that all the failure modes are correctly identified and perfectly cured when they first occur, the overall failure rate must decline faster than the upper bound on the remaining failure rate, $N/(e t)$, where N is the number of correctable failure modes, e is the base of the natural logarithms, and t is time.

The expected progress of reliability growth depends on the number and failure rates of the correctable failures that are originally present. The maximum reliability growth equals the total correctable failure rate. The time to achieve reliability growth depends on the MTBFs of the individual failure modes. The MTBF (Mean Time Before Failure) is the inverse of the failure rate. A failure is more likely to occur at the time near its MTBF than long before or after. For a single correctable failure mode, the expected failure rate decline exactly equals the $1/(e t)$ bound when

test time equals the MTBF. After that time, it declines much more rapidly than $1/t$. A group of failures with similar rates behaves the same way. The clustered failures of a narrow gamma distribution occur near the MTBF and the expected failure rate then declines rapidly. If the failure rates are very different, the total failure rate is much less than the $1/(e t)$ bound. Only the few failure modes with MTBFs close to the current test time contribute substantially to the current failure rate and reliability growth, and the time to reduce the failure rate stretches beyond the longest MTBF corresponding to the lowest failure rate.

C. The heroic reliability effort, duration, and growth metrics

Heroic reliability growth requires correctly diagnosing all the correctable failure modes and then removing the failures without introducing any new ones. If this is done, the failure rate will decline faster than the bound $1/(et)$. For this bound, the abcd reliability growth model $a t^{-b}$ would equal $(1/e) t^{-1}$, so the worst case bound reliability growth exponent for a heroic effort is $b = 1$. If no effort is made to reduce the failure rate, the reliability growth time exponent is 0 and the failure rate is constant. It seems appropriate to define the reliability growth exponent “b” as the *heroic reliability growth effort metric*.

The heroic reliability effort metric varies from 0 for no effort to 1 and even higher if the failure reduction problem is unusually easy. The $1/(et)$ is an upper bound on the failure rate during the reliability growth effort. The failure rate can fall faster if, for instance, there are only a few failure modes with high failure rates. However, there are many more likely reasons why the failure rate decline, which equals the heroic reliability effort metric, can be much less than 1, even approaching 0. If a failure is not detected or corrected the first time it occurs, failure rate reduction is slowed. A particular failure mode may never be removed, preventing reliability growth. An attempted fix may introduce a new failure mode, effectively restarting the reliability growth clock.

Measured reliability growth provides another metric. In the abcd reliability growth model, with a failure rate $= a t^{-b} + c$, the minimum final failure rate is “c,” the residual random failure rate. But all correctable failures are removed only if the failure rate reduction process continues without end. If the process is terminated at t_d , the final failure rate is “c + d,” where $d = a t_d^{-b}$. The failure rate reduction achieved is $\text{Max} - (c+d)$, where Max is the highest initial failure rate rate, $\text{Max} = a t_1^{-b} + c$, and t_1 is the time of the first failure. The greatest possible failure rate reduction, achieved over impossibly long time, is $\text{Max} - c$. It seems appropriate to define achieved failure rate reduction ratio $(\text{Max} - c - d)/(\text{Max} - c)$ as the *heroic reliability improvement metric*. Computing this metric requires that the model distinguish the remaining correctable failure rate, “d,” from the uncorrectable continuous failure rate, “c.”

The abcd model assumes that the failure reduction effort stops at some time t_d . At this point in time, the remaining correctable failure rate is d. With the model’s continuing unending decline in the probability of failure, there is never a time when all correctable faults have been fixed. Suppose that the objective is to reduce the correctable failure rate to r, and that this requires time t_r . From the reliability growth model, $r = a t_r^{-b}$, and time $t_r = \text{EXP}[-\text{LN}(r/a)/b]$. As a difficult reference target, the remaining correctable failure rate, r, is set equal to $r = 0.1 c$, ten percent of the originally expected random failure rate. The final failure rate at time t_r would then be $a t_r^{-b} + c = 0.1 c + c = 1.1 c$. If the reliability growth effort stops at t_d but should extend to t_r , it seems appropriate to define the *heroic reliability growth duration metric* as t_d/t_r . For $r = c$, $t_r = \text{EXP}[-\text{LN}(0.1 c/a)/b]$.

A combined metric, the *heroic reliability growth metric*, can be created by multiplying the three reliability growth metrics for effort, improvement, and duration.

VI. Reliability growth analysis of the Duane-Crow data set

The “abcd” model is applied to the Duane-Crow reliability growth data set. The “abcd” parameters are derived and the heroic metrics computed. These are shown in Table 1.

Table 1. Duane-Crow reliability growth model parameters and metrics.

abcd model parameters						metrics				gamma model	
						effort	improvement	duration	reliability growth		
a	b	c	d	t_d	t_r	b	$(\text{Max} - c - d) / (\text{Max} - c)$	t_d/t_r		alpha	beta
0.97	0.83	0.14	0.01	200	175.5	0.83	0.99	1.14	0.94	0.080	0.938

Figure 3 showed a rough fit to the Duane-Crow data set, with failure rate equal to $1.11 t^{-0.5}$ to time 100. A better fit is computed mathematically using the abcd model.

Failure rate = $a t^{-b} + c = 0.97 t^{-0.83} + 0.14$ from $t = 0$ until the time $t_d = 200$ when $0.97 \cdot 100^{-0.83} = d = 0.01$.
 $= c + d = 0.14 + 0.01 = 0.15$ after time 200.

Figure 11 shows the Duane-Crow failure rate data, the abcd model, and the $N/(e t)$ bound.




Figure 11. The Duane-Crow failure rate data, abcd model, and $N/(e t)$ bound.

Figure 11 shows all the data out to time = 400 and the data and model include the constant failure rate $c = 0.14$. The abcd model parameters were computed using the $N(t)/T$ data out to time = 200 without the constant failure rate c , which was added back for Figure 11. The two-part model combines initial reliability growth with a constant background failure rate and fits the data closely. The upper bound on the failure rate, $N/(e t)$, applies only to the correctable failures and so falls below the constant failure rate. The bound is correct, as shown in Figure 12, which does not include the constant failure rate.




Figure 12. The Duane-Crow failure rate data, models, and bound without the constant failure rate $c = 0.14$.

In Figure 12, the blue dots are the $N(t)/t$ data minus $c = 0.14$ and the red line is the abcd model fit to that data out to time = 200. Figure 13 shows that this model fits the data out to time = 400. The broken orange line is the abcd model fit using data only out to time = 100. For time 0 to 100, the reliability effort metric, b , was higher, 0.99, than the 0.83 found for the full period from 0 to 200. The failure rate decline is rapid and both abcd models quickly converge to the data.

A gamma distribution based estimate of the reliability growth was also developed and is shown as a green line. The $N(t)$ data indicates the specific times of the failures. The failure time can be taken as a rough estimate of the MTBF of the occurring failure mode, and the failure rate is the inverse of the MTBF. The average value and the standard deviation of the estimated failure rates is computed and used to calculate the α and β of the gamma distribution. The gamma distribution of the correctable failure rates provides an estimate of the achievable failure rate decline.

$$\lambda_{cor}(t) = \alpha \beta N / (1 + \beta t)^{-(\alpha+1)}$$

For the Duane-Crow data, the α is quite low, 0.080, indicating that the distribution of failure rates is peaked at the lower rates and that the decrease in the failure rate is relatively slow. The equation for $\lambda_{cor}(t)$ shows that it decreases as $t^{-1.080}$. The gamma distribution estimate also converges quickly to the data.

The $N/(e t)$ bound is an upper bound that is significantly above the actual correctable failure rate data. It is about ten times higher than the data from time = 200 to 400.

Significant reliability growth was achieved, with a strong continued effort. The model predicts that continuing the reliability growth effort at the same level from time 200 to 400 would have reduced d from 0.012 to 0.007 and $c + d$ from 0.147 to 0.142, which would be a very small gain for doubling the reliability growth time. The heroic reliability duration metric would have increased from 0.92 to 0.96. The heroic reliability improvement metric would have increased from 0.988 to 0.993. The Duane-Crow data set does show dramatic reliability growth until time 100, with a heroic reliability effort metric of 0.985. There is a small but observable reliability growth from time 100 until time 200, but from time 200 to 400, the failure rate decrease is less than one-tenth of the constant random failure rate.

VII. Manned space systems reliability growth

Failure rate data and the possibility of reliability growth are considered for the space shuttle, the International Space Station (ISS), and the ISS Carbon Dioxide Removal Assembly (CDRA). The space shuttle demonstrated significant reliability growth. The shuttle was extensively maintained, refurbished, and upgraded after each flight. The ISS failure rate has been approximately constant over time. The ISS CDRA reliability has been roughly constant. Reliability growth requires an active program to discover and analyze failures and make improvements to

fix the discovered design deficiencies. The space shuttle program achieved reliability growth. Failures on ISS and with the CDRA have not decreased.

A. Space shuttle

Figure 13 shows the space shuttle failure rate data, the abcd model plot, and a $k/(e t)$ bound.




Figure 13. The space shuttle failure rate data, the abcd model plot, and a $k/(e t)$ bound.

The cumulative failure rate $N(t)/t$ data is for the space shuttle Columbia for its first 10,000 operating hours. The data was reported as $N(t)/t$ rather than failure times. The other shuttles showed similar failure rate declines. (Shishko, SP-6105) The increase in failure rate at about 4,000 operating hours follows the interruption of shuttle flights after the 1986 Challenger accident.

The shuttle was extensively refurbished, maintained, and upgraded after each flight. A large staff had the time it needed in ideal hangar conditions, which is a significant contrast to the ISS which has been in continual orbital flight. The reduction in failure rate for the space shuttle proves the possibility of reliability growth in a space system.

The two-part abcd model combines continuing reliability growth with a constant background failure rate and fits the data closely. The bound that was fitted to the initial failure rate, $k/(e t) + 0.05$, tracks the corrected failures plus the constant failure rate. The “abcd” parameters and the heroic metrics are shown in Table 2.

Table 2. Manned space reliability growth model parameters and metrics.

Mission	abcd model parameters						metrics			
							effort	improvement	duration	reliability growth
	a	b	c	d	t_d	t_r	b	$(\text{Max} - c - d) / (\text{Max} - c)$	t_d/t_r	
Space shuttle	19.98	0.70	0.05	0.03	9,700	139,694	0.70	0.97	0.07	0.05
ISS	4.26	1.57	11.00	1.01	2.5	2.4	1.57	0.92	1.06	1.53
CDRA	1.43	0.125	0.50	2.10	11	4.48E+11	0.125	-0.05	0.00	0.00

For the space shuttle, the reliability effort, the parameter “b,” was a high 0.70. Until 3,700 on-orbit hours, before Challenger, it was 0.82, significantly higher. The failure rate improvement was substantial, 0.97, from an initial normalized failure rate of 1.0 down to $d = 0.03$, discounting the constant failure rate of $c = 0.05$. Failure reduction efforts continued throughout the entire 9,700 hours recorded in the data. However, the slowing reliability growth,

and the ambitious final target for remaining correctable failure rate of $0.1 \text{ c} = 0.005$, result in a very long time $t_r = 139,694$ hours to reach the final target. Thus the duration and reliability growth metrics are low.

Even though the space shuttle demonstrated significant reliability growth, the data and model suggest that more may have been done. The initial effort metric of 0.82 was higher than the long term one of 0.70. the remaining correctable failure rate $d = 0.03$ was 60% of the modeled continuing random failure rate of $c = 0.05$. A few early failures can greatly increase the $\text{Max } N(t)/t$ and so can exaggerate reliability growth. This suggests that long term continued detection, analysis, and removal of failure modes is more important in reliability growth. And, just as an unduly large number of initial failures can exaggerate reliability growth, so can the accepting a relatively high continuing uncorrectable failure rate.

Down time can reduce readiness, which may explain the increase in failure rate at the return to flight after the Challenger tragedy. One of the many lessons of Challenger is that there can be two different approaches to reliability. Failure events can be treated in two different ways, as ordinary and expected or as extraordinary and unexpected. The process of conducting a test to discover and remove failure modes is ordinary and expected. The circumstances that produce the failures are intentional, ordinary, and expected. The design improvements are intentional, ordinary, and expected. But the unusual cold that disabled the Challenger O-rings and the solid rocket booster explosion that followed were extraordinary and unexpected. Some knew of the problem and strongly suggested urgent intervention but, since there could be no proof that a failure would occur, management was able to discount the possibility and launch into tragedy.

There is a crucial difference between fixing the failures that do occur and fixing the potential failures that might occur. It is easier to do what is obviously necessary than fix something that may not matter. The essence of a high reliability culture is not just removing all occurring failure modes to achieve high everyday operating reliability, it consists of pursuing all anomalies, understanding any off-nominal occurrences, and doing whatever is needed to eradicate any lurking potential problems. It requires heroic effort to achieve significant reliability growth, but it requires hypersensitive alertness and deep paranoid fear to achieve the highest possible safety.

Avoiding surprising hazards is as important as achieving reliability growth. Consider life support. It includes continually operating systems that provide oxygen and water vital to the crew and must be highly reliable. Achieving achieving reliability growth requires continued testing. Life support also includes suppressing fire and preventing atmosphere loss, which hazards have caused loss of crew. Achieving safety requires testing emergency response functions.

B. International Space Station (ISS)

Figure 14 shows the ISS failure rate data, the abcd model plot, and a $k/(e \text{ } t)$ bound.




Figure 14. The ISS failure rate data, the abcd model plot, and a $k/(e \text{ } t)$ bound.

The cumulative failure rate $N(t)/t$ data for the ISS is based on the number of unscheduled maintenance actions from February 1999 to February 2011. (Cirillo et al., 2011) The number of pressurized modules grew from two to fourteen over this period, and the data is normalized as failure rate per year per pressurized module. Since the ISS

was not permanently inhabited for the first two years, the number of unscheduled maintenance actions was much lower.

The abcd model fits the data closely only for the first two and one-half years when the ISS was mostly uninhabited. The bound that was fitted to the initial total failure rate, $k/(e t) + 11$ tracks the abcd model well. The “abcd” parameters and the heroic metrics are shown in Table 2. For the ISS, the initial reliability effort, the parameter “b,” was a very high 1.57 for the first two and one-half years, but it then became 0. The overall failure rate dropped from 24 per year to 12 after two and one-half years, but this included the constant rate of $c = 11$. The correctable failure rate improvement was substantial, 0.92, from the initial failure rate of 24 less the constant failure rate of $c = 11$, down to $d = 1$. The demonstrable failure reduction efforts continued only two and one-half years, but this was more than sufficient to achieve the final target for remaining correctable failure rate of $0.1 c = 1.1$. The elapsed duration was sufficient to achieve the target reliability, so the duration metric is greater than 1. The overall reliability growth metric is thus very high, 1.53.

Regardless of the good model fit to the early data, and regardless of the favorable reliability growth metrics, the ISS failure data do not show significant reliability growth. The ISS measured failure rate, $N(t)/t$, decreased only by one-third over the first twelve years, from 24 to 16. To achieve reliability growth, failures must be analyzed and the design changed to prevent the problems recurring. The ISS is a very difficult operational environment. Developing and implementing improvements requires a long time cycle, even compared to its long operational life.

Russell and Klaus found that Environmental Control and Life Support (ECLS) repair was a major component of ISS maintenance. They also found that the ISS ECLS maintenance load was constant over an 865 day period. (Russell and Klaus, 2006) (Russell et al., 2006)

There are several reasons that the ISS failure data do not show significant reliability growth. First, the reliability growth effort was relatively brief compared to the later operating time. Second, with the ISS not inhabited, fewer failures would occur and fewer unscheduled maintenance actions, the data recorded here, would be possible. Third, the very high continuing and apparently accepted failure rate indicate that reliability growth was not able to have high enough high priority to overcome its intrinsic difficulties.

C. Carbon Dioxide Removal system (CDRA)

Figure 15 shows the ISS Carbon Dioxide Removal Assembly (CDRA) failure rate data, a constant approximation, and the abcd model.




Figure 15. The CDRA failure rate data, a constant approximation, and the abcd model.

The ISS CDRA has been operating on board ISS since 2001. The number of failures per year was determined from yearly ISS life support system status papers. (Reuter, 2000-01-2248, Reuter and Reysa, 2001-01-2386, Gentry et al., 2002-01-2495, Williams et al., 2003-01-2589, Williams and Gentry, 2004-01-2382, Williams and Gentry, 2005-01-2777, Williams and Gentry, 2006-01-2055, Williams and Gentry, 2007-01-3098, Williams and Gentry, 2008-01-2131, Williams and Gentry, 2009-01-2415, Williams et al., 2010-6180, Williams et al., 2012-3612) Two units were on board during the last year reported and the cumulative failure rate is given per unit.

The data was gathered in the form of the number of failures per year. No failures were reported in years 4, 5, and 6, which were 2004, 2005, and 2006. This gives an impression of reliability growth, but many failures occurred in the following years.

A reasonably good fit to the data is a constant 1.76 failures per year. The abcd model is not much different. The “abcd” parameters and the heroic metrics are shown in Table 2. For the CDRA, the reliability effort parameter “b,” was only 0.125, not much different from 0. The initial and final failure rate was 2 per year, no real improvement, and the abcd model suggests a slight loss of reliability. The slow rate of improvement indicates a fantastically long time to achieve the final reliability target. The reliability growth metric is 0.

D. Oxygen Generation Assembly (OGA) reliability growth

Figure 16 shows the ISS Oxygen Generation Assembly (OGA) $N(t)/t$ failure rate data and three abcd models based on different periods of the data. Different abcd models are reflect changing behavior over time.




Figure 16. The OGA failure rate data and three abcd models.

Two papers summarize the failure and maintenance events of the OGA now on the ISS. (Takada et al., 2015-115 (Bagdigian et al., 2015-094). The operational problems reported in these papers were consolidated to provide the $N(t)/t$ failure rate data shown in Figure 16.

Three abcd models were fit to the data for three increasing time periods from initial operation out to 0.5, 1.4, and the full 6.2 years. The “abcd” parameters and the heroic metrics are shown in Table 3 for the three abcd models.

Table 3. OGA growth model parameters and metrics.

Data span, years	abcd model parameters						metrics			
							effort	improvement	duration	reliability growth
	a	b	c	d	t_d	t_r	b	$\frac{(\text{Max} - c - d)}{(\text{Max} - c)}$	t_d/t_r	
0 - 0.5	3.17	0.85	5.00	5.44	0.53	8.78	0.85	0.91	0.06	0.05
0 - 1.4	8.59	0.40	5.00	7.53	1.39	1,223	0.40	0.88	0.00	0.00
0 - 6.2	14.44	0.02	5.00	13.95	6.24	7.47E+76	0.02	0.77	0.00	0.00

The initial abcd model fits the data closely for the first half year. The initial reliability effort, the parameter “b,” was an effective 0.85. Significant reliability improvement, 0.91, was achieved. However, this initial effort ceased at time $t_d = 0.53$ with a large remaining correctable failure rate, $d = 5.55$. Reducing this to 0.1 $c = 0.5$ would require 8.78 years, so the duration achieved is low, 0.06, and the reliability growth metric is low, 0.05.

After $t_d = 0.53$, the failure rate quickly increases and then declines much more slowly. The second abcd model fits the data closely for the initial 1.4 years. The reliability effort, the parameter “b,” drops to 0.40. The higher current failure rate reduces the achieved reliability improvement. Eliminating the new higher correctable failure rate with reduced effort would take impossibly long, and the duration and reliability growth metrics drop to 0.

After time = 1.4 years, there is a large step increase in the failure rate followed by a continuing much slower failure rate reduction. The entire data set is used to calibrate the third abcd model, which is roughly a constant 19.44 failures per year. The continuing failure rate decreases are offset by the increases at 0.4 and 1.4 years. The reliability effort, the parameter “b,” drops to 0.02. The higher current failure rate reduces the achieved reliability improvement. The time to eliminate the correctable failure rate becomes longer than the universe will exist. The duration and reliability growth metrics are 0. The initial rapid reliability growth could not be sustained and a continuing high failure rate persists. Like the entire ISS and the CDRA, the OGA does not show significant reliability growth.

VIII. Conclusion

The general process of heroic reliability improvement is easy to explain but difficult to implement. If high rate failure modes are diagnosed and then removed by redesign, they will no longer occur and the total failure rate will be reduced. This is easier said than done.

The general process of heroic reliability improvement can be described with simple mathematics. Suppose there are N correctable high probability failure modes. Since they have high failure rate, they will all probably occur early in testing, say before time T . After testing starts, the number of failures will gradually increase until N have occurred at time T . The failure rate is then $N(t)/t = N/T$. If all of the N correctable failure modes have been removed, no more correctable failures will occur. As time goes on, the cumulative failure rate $N(t)/t = N/t$ will decline as $1/t$ or t^{-1} . This is the best case of failure rate decline. The commonly used Duane-Crow reliability growth model stops at this point, with the main issue being the time exponent of failure rate decline.

A practical complication, not considered in the heroic model, is that all the correctable failure modes may not be removed. Suppose none of the N are removed. All will reoccur. Their average time of initial occurrence, their MTBF, is $T/2$. Their average failure rate is the inverse of the MTBF, $\lambda = 1/\text{MTBF} = 2/T$. The total failure rate is a constant $N\lambda = 2N/T$. Failures that are correctable but not corrected continue to fail at a constant rate.

The usual reliability model is of a system with many components that each have a constant acceptable failure rate. The system failure rate is the sum of the constant component failure rates. This can be regarded as a best case failure rate. Design mistakes, bad components, and unexpected conditions may cause a much greater initial failure rate, with most of the problems treated as correctable failure modes.

A problem in using the Duane-Crow reliability growth model is that it does not include an acceptable residual failure rate due to uncorrectable failure modes. The assumption is that reliability growth will continue and the failure rate decrease indefinitely. As the constant rate acceptable failures accumulate beyond the period of reliability growth, the reliability growth time exponent decreases toward zero.

The approach taken here was to combine the heroic failure reduction model with the constant acceptable failure rate model, to form the abcd model. The failure rate $N(t)/t = a t^{-b} + c$, where $a t^{-b}$ describes heroic reliability growth and c is the acceptable failure rate. The parameter d represents the additional constant failure rate due to correctable but uncorrected failure modes.

The abcd model seems to be useful in describing and understanding reliability growth data. It seems to be able to predict future reliability growth, but only if the current reliability effort continues at the same level and there are no unforeseen surprises.

As a practical matter, the great difficulty in analyzing failures and implementing redesigns on the ISS makes it very difficult to improve ISS reliability. ISS and ISS life support have not shown heroic reliability growth. The reliability of the ISS CDRA and OGA is not significantly improving. “It continues to be a challenge to provide the functionality necessary to support the six crew members on ISS with all of the problems that have occurred with the new regenerative (life support) and CDRA.” (Williams et al., 2012-3612) The analysis of heroic reliability growth suggests that longer testing with diligent failure mode correction is required to reduce actual failure rates to more acceptable levels. It would seem easier to do most of this on the ground rather than on-orbit.

The most useful application of the heroic reliability improvement model may be in planning and guiding the failure rate reduction process. The expected system failure rate is usually computed during design. Some estimate of the likely number and failure rates of unplanned correctable failures, based on similar past experience, could be used to estimate the required test time and level of effort needed to remove failure modes. As the reliability improvement process is carried out, the current failure rate and updated model could be used to track progress and adjust planning.

References

- Bagdigian, R. M., Dake, J., Gentry, G., and Gault, M., "International Space Station Environmental Control and Life Support System Mass and Crewtime Utilization In Comparison to a Long Duration Human Space Exploration Mission," ICES-2015-094, 45th International Conference on Environmental Systems 12-16 July 2015, Bellevue, Washington.
- Bishop, P., and Bloomfield, R., "A Conservative Theory for Long-Term Reliability-Growth Prediction, IEEE Transactions on Reliability, vol. 45, no. 4, 1996.
- Cirillo, W., Goodliff, K., Aaseng, G., Stromgren, C., and Andrew Maxwell, A., "Supportability for Beyond Low Earth Orbit Missions," AIAA 2011-7231, AIAA SPACE 2011 Conference & Exposition, 27 - 29 September 2011, Long Beach, California.
- Gentry, G. J., Reysa, R. P., and Lewis, J. F., "International Space Station (ISS) Environmental Control and Life Support (ECLS) System Equipment Failures, Causes, and Solutions February 2001 - February 2002." SAE 2002 - 01 - 2495; 32nd International Conference on Environmental Systems; San Antonio, Texas, July 2002.
- Hansen, R. C., Overall Equipment Effectiveness, Industrial Press Inc, New York, New York, 2001.
- Jones, H. W., "Common Cause Failures and Ultra Reliability," AIAA 2012-3602, 42nd International Conference on Environmental Systems, 15 - 19 July 2012, San Diego, California.
- Jones, H. W., "Reliability Growth in Space Life Support Systems," ICES 2014-075, 44th International Conference on Environmental Systems, 13-17 July 2014, Tucson, Arizona.
- MIL-HDBK-189C, Department of Defense Handbook Reliability Growth Management, 14 June 2011.
- Reuter, J. L., "International Space Station Environmental Control and Life Support Status: 1999 – 2000," SAE 2000-01-2248, 30th International Conference on Environmental Systems; Toulouse, France, July 2000.
- Reuter, J. L., and Reysa, R. P., "International Space Station Environmental Control and Life Support Status: 2000 – 2001," SAE 2001-01-2386; 31st International Conference on Environmental Systems; Orlando, Florida, July 2001.
- Russell, J. F. and Klaus, D. M., "Maintenance, Reliability and Policies for Orbital Space Station Life Support Systems," Reliability Engineering and System Safety, Vol. 92, No. 6, 2006, pp. 808-820.
- Russell, J. F. and Klaus, D. M., "Maintenance, Reliability and Policies for Orbital Space Station Life Support Systems," Reliability Engineering and System Safety, Vol. 92, No. 6, 2006, pp. 808-820.
- Russell, J. F., Klaus, D. M., Mosher, T., "Applying analysis of international space station crew-time utilization to mission design," J Spacecraft Rockets, 2006;43(1):130–6.
- Russell, J. F., Klaus, D. M., Mosher, T., "Applying analysis of international space station crew-time utilization to mission design," J Spacecraft Rockets, 2006;43(1):130–6.
- Shishko, R., NASA Systems Engineering Handbook, NASA-SP-6105, June 1995.
- Takada, K. C., Ghariani, A. E., and Van Keuren, S., "Advancing the Oxygen Generation Assembly Design to Increase Reliability and Reduce Costs for a Future Long Duration Mission," ICES-2015-115, 45th International Conference on Environmental Systems, 12-16 July 2015, Bellevue, Washington.
- Williams, D. E., and Gentry, G., "International Space Station Environmental Control and Life Support System Status: 2003 – 2004," SAE 2004-01-2382; 34th International Conference on Environmental Systems; Colorado Springs, Colorado, July 2004.
- Williams, D. E., and Gentry, G., "International Space Station Environmental Control and Life Support System Status: 2004 – 2005," SAE 2005-01-2777; 35th International Conference on Environmental Systems; Rome, Italy, July 2005.
- Williams, D. E., and Gentry, G., "International Space Station Environmental Control and Life Support System Status: 2005 – 2006," SAE 2006-01-2055; 36th International Conference on Environmental Systems; Norfolk, Virginia, July 2006.
- Williams, D. E., and Gentry, G., "International Space Station Environmental Control and Life Support System Status: 2006 – 2007," SAE 2007-01-3098; 37th International Conference on Environmental Systems; Chicago, Illinois, July 2007.
- Williams, D. E., and Gentry, G., "International Space Station Environmental Control and Life Support System Status: 2007 – 2008," SAE 2008-01-2131; 38th International Conference on Environmental Systems; San Francisco, California, July 2008.
- Williams, D. E., and Gentry, G., "International Space Station Environmental Control and Life Support System Status: 2008 – 2009," SAE 2009-01-2415, 39th International Conference on Environmental Systems; Savannah, Georgia, July 2009.
- Williams, D. E., Dake, J., and Gentry, G., "International Space Station Environmental Control and Life Support System Status: 2009 – 2010," AIAA 2010-6180, 40th International Conference on Environmental Systems; Barcelona, Spain, July 2010.
- Williams, D. E., Dake, J., and Gentry, G., "International Space Station Environmental Control and Life Support System Status for the Prior Year: 2010-2011, AIAA 2012-3612, 42nd International Conference on Environmental Systems, 15 - 19 July 2012, San Diego, California.
- Williams, D. E., Lewis, J. F., and Gentry, G., "International Space Station Environmental Control and Life Support Status: 2002 – 2003," SAE 2003-01-2589; 33rd International Conference on Environmental Systems; Vancouver, British Columbia, Canada, July 2003.
- Yamada, S., and Osaki, S., "Reliability growth models for hardware and software systems based on nonhomogeneous Poisson processes: a survey," Microelectronics Reliability, Vol. 23, No. 1, pp. 91-112, 1983.